

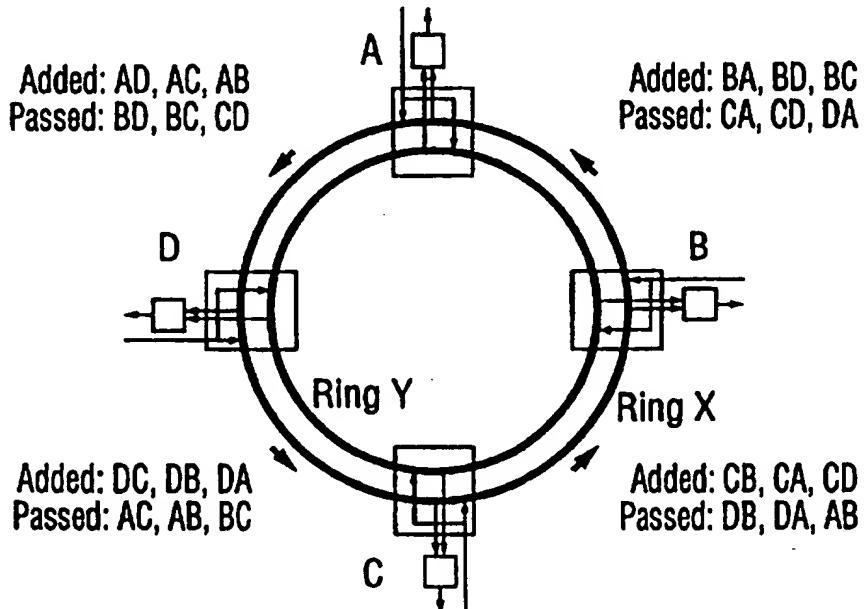


INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ :	A1	(11) International Publication Number:	WO 97/27718
H04Q 11/04, H04L 12/43		(43) International Publication Date:	31 July 1997 (31.07.97)

(21) International Application Number:	PCT/CA97/00052	(81) Designated States: AL, AM, AT, AU, AZ, BB, BG, BR, BY, CA, CH, CN, CZ, DE, DK, EE, ES, FI, GB, GE, HU, IL, IS, JP, KE, KG, KP, KR, KZ, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, TJ, TM, TR, TT, UA, UG, US, UZ, VN, ARIPO patent (KE, LS, MW, SD, SZ, UG), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, ML, MR, NE, SN, TD, TG).
(22) International Filing Date:	27 January 1997 (27.01.97)	
(30) Priority Data:		
9601692.8	27 January 1996 (27.01.96)	GB
(71) Applicant (for all designated States except US):	NEWBRIDGE NETWORKS CORPORATION [CA/CA]; 600 March Road, P.O. Box 13600, Kanata, Ontario K2K 2E6 (CA).	
(72) Inventor; and		Published
(75) Inventor/Applicant (for US only):	MITCHELL, Charles [CA/CA]; 78 Summerwalk Place, Nepean, Ontario K2G 5Y4 (CA).	With international search report. Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.
(74) Agent:	MITCHELL, Richard, J.; Marks & Clerk, Station B, P.O. Box 957, Ottawa, Ontario K1P 5S7 (CA).	

(54) Title: ATM NETWORK WITH RING ARCHITECTURE



(57) Abstract

A communications network comprises a plurality of nodes interconnected by a pair of packet-based rings, for example ATM rings. At each said node there is a ring interface comprising a device for extracting packets from either ring destined for that node, a device for adding packets from that node to either ring, and a device for passing directly through packets on each ring not destined for that node. This arrangement establishes a virtual mesh structure with offers good fault recovery without unnecessary wastage of bandwidth.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AM	Armenia	GB	United Kingdom	MW	Malawi
AT	Austria	GE	Georgia	MX	Mexico
AU	Australia	GN	Guinea	NE	Niger
BB	Barbados	GR	Greece	NL	Netherlands
BE	Belgium	HU	Hungary	NO	Norway
BF	Burkina Faso	IE	Ireland	NZ	New Zealand
BG	Bulgaria	IT	Italy	PL	Poland
BJ	Benin	JP	Japan	PT	Portugal
BR	Brazil	KE	Kenya	RO	Romania
BY	Belarus	KG	Kyrgyzstan	RU	Russian Federation
CA	Canada	KP	Democratic People's Republic of Korea	SD	Sudan
CF	Central African Republic	KR	Republic of Korea	SE	Sweden
CG	Congo	KZ	Kazakhstan	SG	Singapore
CH	Switzerland	LI	Liechtenstein	SI	Slovenia
CI	Côte d'Ivoire	LK	Sri Lanka	SK	Slovakia
CM	Cameroon	LR	Liberia	SN	Senegal
CN	China	LT	Lithuania	SZ	Swaziland
CS	Czechoslovakia	LU	Luxembourg	TD	Chad
CZ	Czech Republic	LV	Latvia	TG	Togo
DE	Germany	MC	Monaco	TJ	Tajikistan
DK	Denmark	MD	Republic of Moldova	TT	Trinidad and Tobago
EE	Estonia	MG	Madagascar	UA	Ukraine
ES	Spain	ML	Mali	UG	Uganda
FI	Finland	MN	Mongolia	US	United States of America
FR	France	MR	Mauritania	UZ	Uzbekistan
GA	Gabon			VN	Viet Nam

- 1 -

ATM NETWORK WITH RING ARCHITECTURE

This invention relates to packet-based networks, such as ATM (Asynchronous Transfer Mode) networks, having a ring architecture. The invention is applicable to both local and wide area networks.

There is a need within networks to ensure that network operation is not interrupted in the event of infrastructure failures, such as cable cuts and/or equipment failure. The current state of the art is to provide protection with SONET Ring architectures. Bellcore standard GR-1230-CORE describes a Bi-direction Line Switched Line Switched Ring and GR-1400-CORE describes a Uni-directional Path Switched Ring.

In the case of ATM networks, the current state of the art is to send ATM payloads over SONET ring structures. As this is an expensive solution, there is a need to provide a simpler mechanism for ATM network protection. While it is possible to map SONET Ring structures directly into an ATM implementation using the ATM Switch to perform the necessary bridging and re-configuration, the result is an inefficient use of the ATM Switch and a slow response time.

An object of the invention is to provide a wide area network architecture that alleviates the afore-mentioned problems of the prior art.

According to the present invention there is provided a communications network a communications network comprising a plurality of nodes interconnected in a ring architecture. A pair of counter-rotating packet-based rings interconnect the nodes to establish virtual connections therebetween. Each node includes a ring interface and a packet switch connected thereto. The ring interface comprises means for extracting packets from either ring destined for the packet switch, means for adding outgoing packets to the packet switch to either ring, and means for passing directly through packets on each ring not destined for the node. This results in a virtual mesh structure being established between said nodes.

The rings can either be unidirectional or bi-directional. Of course, it will be understood by one skilled in the art that additional rings can be provided without detracting from the advantages offered by the invention.

- 2 -

The described ring structure allows bandwidth to be efficiently shared between the nodes on the ring without the wastage inherent in a hard-wired mesh and without using up valuable ATM switch resources since onward cells are passed directly through the interface without being forwarded to the ATM switch.

The invention also provides a method of establishing communication over a network wherein a plurality of nodes are interconnected in a ring architecture. A pair of counter-rotating packet-based rings are established to interconnect the nodes and provide virtual connections therebetween. At each node packets destined for the node are extracted from the rings, outgoing packets from the node are added to the rings, and packets on each ring not destined for that node are passed directly through. This results a virtual mesh structure being established between the nodes.

The invention will now be described in more detail, by way of example only, with reference to the accompanying drawings, in which:-

Figure 1 shows a four node mesh network;

Figure 2 shows a four node ring network;

Figure 3 shows a ring interface using OC-12c Interface Cards for a Newbridge Networks 36170 ATM switch;

Figure 4 shows a simple uni-directional ATM Switched Ring (UASR);

Figure 5 shows an example of a double ring fault;

Figure 6 shows another example of a double ring fault;

Figure 7a shows a bi-directional ATM Switched Ring (BASR);

Figure 7b shows a portion of the ring shown in figure 7a;

Figure 8 shows a SONET-ATM ring overlay.

Figure 9 is a block diagram of an OC-12c card; and

Figure 1 shows a four-node mesh network forming a SONET ring. There are six possible bi-directional connections between the nodes. i.e. A-B, A-C, A-D, B-C, B-D and C-D. In a mesh network such as this, six point-to-point links are required to complete the

- 3 -

mesh. If these links are OC-3c interfaces with approximately 150 Mbps capacity, since the bandwidth of these individual links is fixed, no advantage can be taken of the variable bandwidth demands of the nodes. If node A on average requires 50 Mbps to node B, 100 Mbps to node C and only 20 Mbps to node D, the unused capacity to these links is wasted.

While it is possible to route traffic via alternative routes to try and make use of this bandwidth, this is a software intense task that requires extra CAC (Connection Admission Control) calculations and uses up valuable switch resources at these tandem nodes.

In Figure 2, the four nodes A, B, C, D are interconnected by an OC-12c ATM Ring. An OC-12c interface has four times the capacity of a OC-3c link, i.e. approx. 600 Mbps. In this case, each segment of the ring carries six virtual paths, resulting in a logical mesh, i.e. each node has a direct virtual path to every other node. The advantage of this arrangement is that all of the paths are carried on the one link. This means that the nodes using less bandwidth free up the capacity for other nodes to use, and when required nodes can use up to the full bandwidth of the Ring, i.e. node A can send a full 600 Mbps to node C.

Two types of ATM ring can be used: a) A Uni-directional ATM Switched Ring (UASR), and b) A Bi-directional ATM Switched Ring (BASR).

In the UASR, the same traffic is transmitted in opposite directions on two counter-rotating Rings. At each node ingress ATM cells come in on a standard cell relay interface. Cells destined for that node are "dropped" into the switching fabric of the node while the remaining cells are queued with the cells to be "added" by that node and sent out the egress port of the standard cell relay interface.

Figure 3 shows how this ring interface maps onto a Newbridge Networks Corporation 36170 OC-12c card at each node. Rings X and Y pass through OC-12c cards 1 and 2 respectively. These are connected through a Fabric Interface Card (FIC) 3 to a Newbridge Networks Corporation 36170 ATM switching Core 4.

- 4 -

On the ingress side of each ring, the Ingress Cell Controller (ICC) 5 on the OC-12c card determines whether the incoming cells are to be dropped onto the UP_ISL by looking at the cell's VPI field. Each node on the ring is addressed with a range of VPI values. The Fabric Interface Card 3 receives drop traffic from each Ring and has the ability to select cells from either, or both Rings.

This ability is provided by, for example, Newbridge Networks Corporation Diablo ASIC 6, which is a two-input to one output cell switching engine. The inputs of the Diablo 6 are connected to the respective rings X, Y. The "through" cells are queued with cells to be added from the DOWN_ISL (Inter-Shelf Link), again using the Diablo switching ASIC 6 and Egress cell controller 7. Alternatively, all traffic from the ring can be brought up the UP_ISL to the switching core where the cells to be dropped are sent down to their appropriate ISL, and through cells are returned down the OC-12 cards DOWN_ISL along with cells to be added to the ring. In either case, an important feature is that traffic to be dropped at a node can be selected from either or both rings.

The rules for receiving data off the ring are quite different from a time division multiplex ring, such as a SONET ring, for example.

Figure 4 shows a simple UASR. If both the X and Y rings are fault-free, any node on the ring is free to receive data from either ring, i.e. the Diablo switch 6 on the FIC 3 selects only one of its inputs to switch directly to its output. If a fault occurs on one of the rings, then all nodes program the Diablo 6 to receive data off the other fault-free ring. If both rings have faults, then all nodes receive data from both rings by enabling both inputs to the Diablo 4.

The common cases are where both rings have faults, and the rings can recover fully, a cable cut on the span connecting two nodes, and a node going out of service. Other double faults can leave the ring operating in a diminished capacity.

It will be appreciated that with ATM data transmission, cells are only transmitted when there is data to send. Consequently, as long as the identical cells are not being received on the two rings, any ATM cell destined for a particular node can be received. In order to ensure that identical cells are not received at any node, it is important to ensure

- 5 -

that when both rings have faults, any node that detects a fault on its incoming link squelch its corresponding counter-rotating link on the same span.

There are several cases of faults on both rings that need to be considered. The simplest and most common case is that of a cable cut in a span connecting two nodes in the Ring. Figure 5 shows illustrates this case. In this case the nodes A, D surrounding the cut will detect the cut on the incoming ring and unnecessarily squelch the traffic in the counter rotating ring. Unnecessarily, in this case because the cell flow in the counter rotating direction could not have been received by any other node due to the cut.

A more generalized case is shown in Figure 6. In this case a double fault results in segregated islands of links. Each island is capable of full-duplex communication with other members in the island. For example in Figure 6, nodes A through F will enjoy complete communication with each other. Note that if nodes G and H had not squelched their traffic, cells from G and H would be received twice at nodes A through F (in the double Ring fault case, all nodes receive cells from both Rings).

During normal operation, in-band ATM cells or part the SONET can be used to indicate to all nodes on the ring that the ring is fault. In the latter case, any node that has detected a fault on one of its incoming links sets an indication to the next node on the uni-directional ring. Any node receiving such an indication on its incoming link sets an indication to the next node on the uni-directional ring.

The bi-directional ATM Switched Ring (BASR) shown in Figures 7a and 7b is a further refinement of the ATM Ring concept where a small change in approach allows for the use of the "extra" bandwidth on the Ring. Extra bandwidth is freed up on the ring by not broadcasting the same cells on both Rings, but selectively sending Add cells on one of the two Rings. Figures 7a and 7b show how node A can transmit cells on the X Ring to node D and transmit cells on the Y Ring to nodes B and C. As a minimum, this results in less traffic on the rings and therefore less congestion at the queuing points. For this type of ring the receiving nodes are programmed to always receive cells from both rings. In the event of ring failures, it is the transmitting function that alters its behavior. Any ring failure will cause all nodes to change from transmitting selectively on a particular ring, to broadcasting all cells on both rings.

- 6 -

The examples discussed so far have been for OC-12c Rings. OC-12c rings map well into the Newbridge Networks Corporation's 36170 ISL (Inter-Shelf Link) architecture. The particularly attractive feature is that the through traffic on the rings does not go into the switching core and that it takes only one ISL to support an OC-12c Ring.

There are two fundamentally different applications for OC-12c Rings. The first is the "backbone" application where OC-12c Rings are used to connect large switches to form a network backbone. The other is a "collector" application where many smaller nodes feed the ring with smaller bandwidths.

In the backbone case, the number of switches on a ring would normally be limited to four or five nodes. This limits the delay and CDV and it also justifies the size of the add and drop rates, i.e. ISLs (smaller add/drop rates would fall under the collector application).

Additional capacity is readily achieved by adding parallel rings. If the number of nodes on the backbone begins to exceed the optimum number, the parallel rings would be staggered to pick up different nodes.

The collector application by definition requires less add/drop bandwidth and generally picks up traffic from a larger number of smaller nodes. In this case the add/drop bandwidth need not be a full ISL. An OC-12c ring card can be made to fit in a UCS slot. This would allow even Standalone 36170 Shelves to be connect in a ring configuration. An application here may be pick up low bandwidth Frame Relay traffic onto a feeder ring (which would connect to a larger ATM backbone ring).

Transmitting selectively out the two OC-12c cards can be accomplished by:

1) a combination of programming the headers appropriately at all Ingress cards and programming the point-to-point filters to select one or both of the OC-12c card addresses, or

2) a hardware ECC function that identifies and enables/disables groups of connections. In this case Ingress connection programming would not change.

ATM Rings can also be made to work well using OC-3c interface cards. In this case, through traffic on the Ring would have to go through the switching core.

- 7 -

Transmitting out both OC-3c cards is readily accomplished by programming the point-to-point filters. Multicast is also easily handled by programming the same multicast groups.

Ingress selection of which Ring to listen to is accomplished by programming the Diablo switches 6 on the FICs 3 to select the appropriate Add bus (or both in the case of a Ring failure).

Transmitting selectively out the two OC-3c cards is accomplished by a combination of programming the Newbridge headers appropriately at all Ingress cards and programming the point-to-point filters to select one or both of the OC-3c card addresses.

Figure 9 is a block more detailed diagram of an OC-12c card (1, 2) (Figure 3). Cells coming off the ring pass through PHY unit 10 to ICC unit 5. Cells not destined for the node pass through diablo unit 1, FIFO 12 and egress PHY unit 13. Cells destined for the node are extracted by Diablo 3 and passed to UP_ISL link 15.

Cells from the node destined for the ring passes through stealth unit 16 to diablo unit 2. Egress cell controller 7 controls the adding of cells from the node to those passing through.

Given that there exists a significant installed base of SONET rings, it is also important that the ATM ring work well with existing SONET Rings.

Figure 8 shows how an ATM ring can be overlaid on a SONET ring. This is important because ATM interfaces become increasingly difficult to design at the higher SONET rates. Interfaces at OC-12c are feasible and cost effective; however cell processing at the next step common step in the hierarchy, OC-48c, becomes extremely difficult to deal with. IN Figure 8, For example, an OC-12c ATM ring is overlaid on an OC-48 SONET ring. BSS is an ATM ring with ATM VP functionality as per Bellcore GR-2837-CORE.

There will be demands to handle aggregate bandwidth of much greater than that offered by OC-12c interfaces. This can be achieved by running multiple OC-12c ATM rings overlaid on simple OC-48 or OC-192 SONET rings.

- 8 -

The described architecture provides good fault protection while minimizing wasted available bandwidth.

Although preferred embodiments of the invention have been illustrated in the accompanying drawings and described in the foregoing detailed description, it will be understood that the invention is not limited to the embodiments disclosed, but is capable of numerous rearrangements, modifications and substitutions of parts and elements without departing from the spirit of the invention which is defined in the following claims.

Claims:

1. A communications network comprising a plurality of nodes interconnected in a ring architecture, characterized in that a pair of counter-rotating packet-based rings interconnect said nodes to establish virtual connections therebetween, and each said node includes a ring interface and a packet switch connected thereto, said ring interface comprising means for extracting packets from either ring destined therefor and passing them to the packet switch, means for adding outgoing packets from the packet switch to either ring, and means for passing packets directly through on each ring not destined for that node, whereby a virtual mesh structure is established between said nodes.
2. A network as claimed in claim 1, characterized in that said rings are unidirectional and counter-rotating.
3. A network as claimed in claim 1, characterized in that said rings are bi-directional.
4. A network as claimed in claim 3, wherein during normal operation said adding means selectively transmits packets on only one of said bi-directional rings, and in the event of a fault broadcasts packets both said rings.
5. A network as claimed in any of claims 1 to 3 characterized in that said network is an ATM network and said packets are ATM cells.
6. A network as claimed in any of claims 1 to 4 characterized in that said ATM network is overlaid onto a time division multiplexed (TDM) network.
7. A network as claimed in claim 5 wherein said TDM network is a SONET network.
8. A method of establishing communication over a network wherein a plurality of nodes are interconnected in a ring architecture, characterized in that a pair of counter-rotating packet-based rings are established to interconnect said nodes and provide virtual connections therebetween, and that at each node packets destined for the node are extracted from the rings, outgoing packets from the node are added to the rings, and packets on each ring not destined for that node are passed directly through, whereby a virtual mesh structure is established between said nodes.

- 10 -

9. A method as claimed in claim 8, characterized in that said rings are unidirectional and counter-rotating.
10. A method as claimed in claim 9, wherein when both said rings are fault-free said interface selects one of said rings to receive data, when one of said rings has a fault said interface device receives data off the other ring, and when both said rings have faults said interface device receives data of both said rings.
11. A method as claimed in claim 9 or 10 characterized in that upon detection of a fault, said interface device inhibits outgoing data transmission over a span to an adjacent node on which a fault has been detected.
12. A method as claimed in claim 8, characterized in that said rings are bi-directional.
13. A method as claimed in claim 11, wherein during normal operation packets are selectively transmitted on only one of said bi-directional rings, and in the event of a fault packets are broadcast on both said rings.
14. A method as claimed in any of claims 8 to 13, characterized in that said network is an ATM network and said packets are ATM cells.
15. A method as claimed in any of claims 8 to 14 characterized in that said ATM network is overlaid onto a time division multiplexed (TDM) network.
16. A method as claimed in claim 15 characterized in that said TDM network is a SONET network.
17. A method as claimed in 14, characterized in that fault information is carried by in-band ATM cells transmitted on the rings.
18. A method as claimed in 17, characterized in that fault information is carried by SONET overhead.

1/5

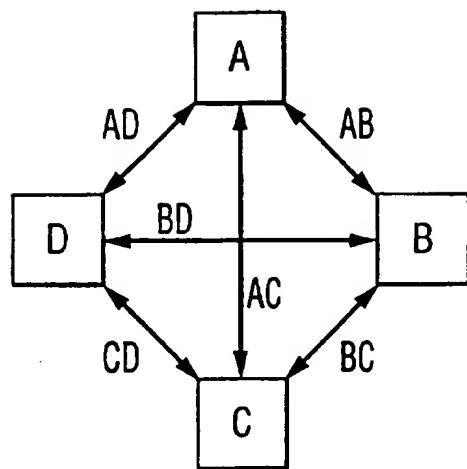


FIG. 1

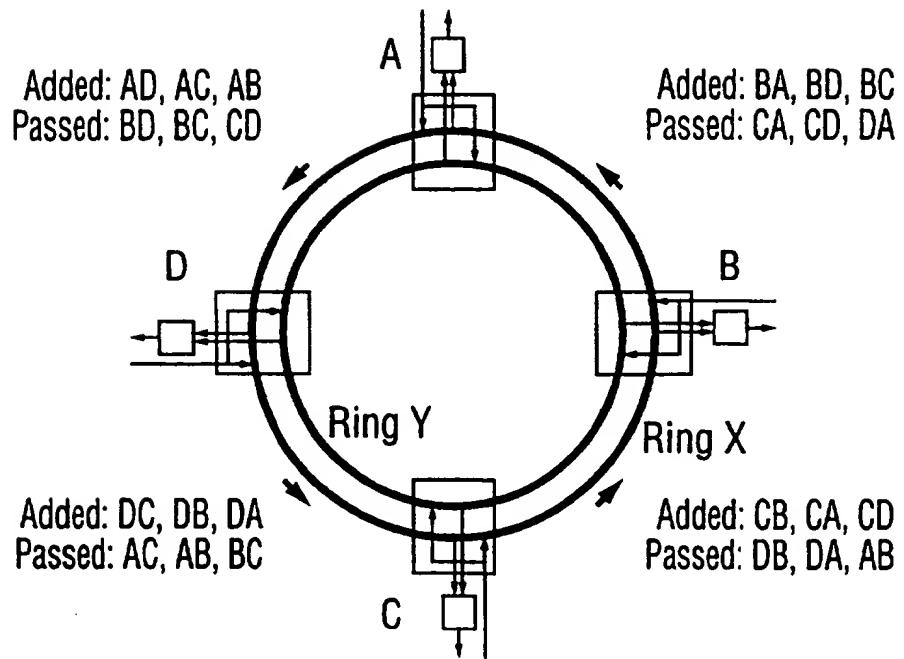
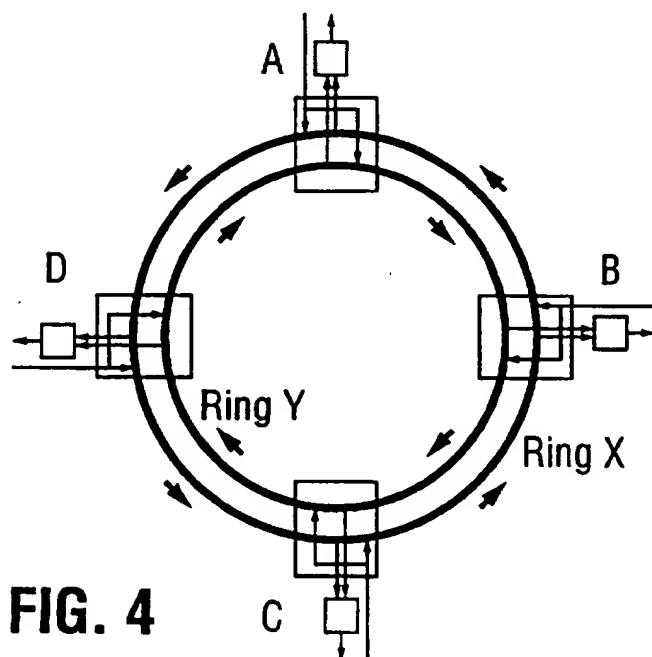
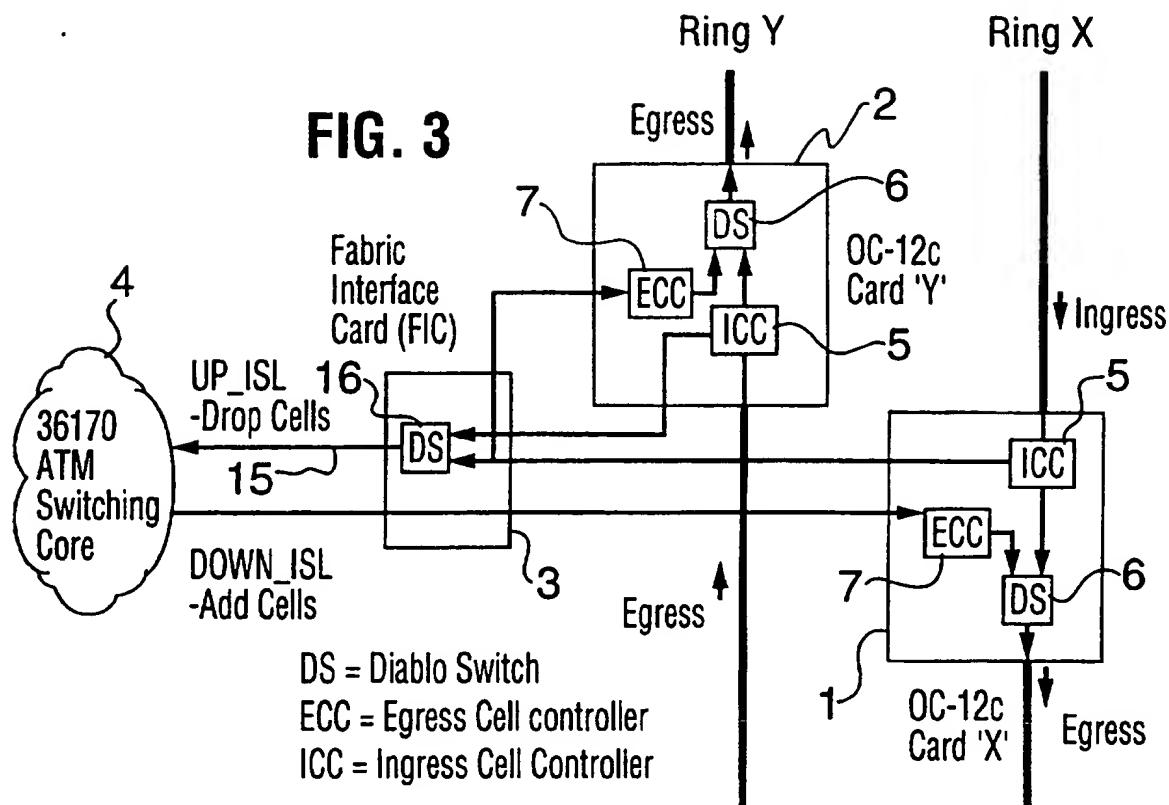


FIG. 2



Simple UASR:

- 1) TX all traffic on both Rings
- 2) If no failure
Rx all traffic from either
If one Ring fails
RX traffic from other Ring
If both Rings fail*
RX traffic from both Rings

*both directions of a span

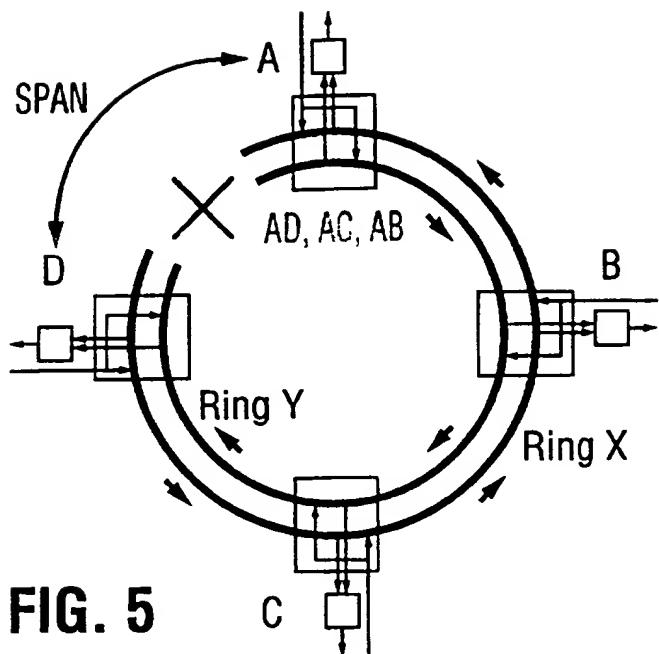


FIG. 5

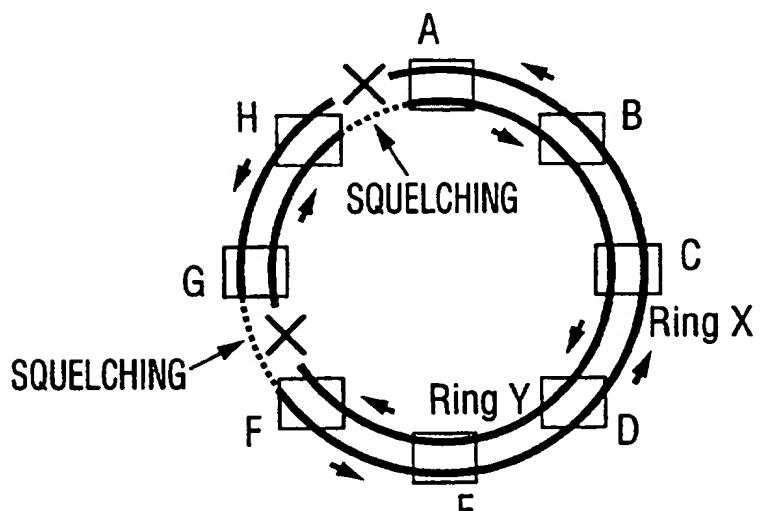


FIG. 6

4/5

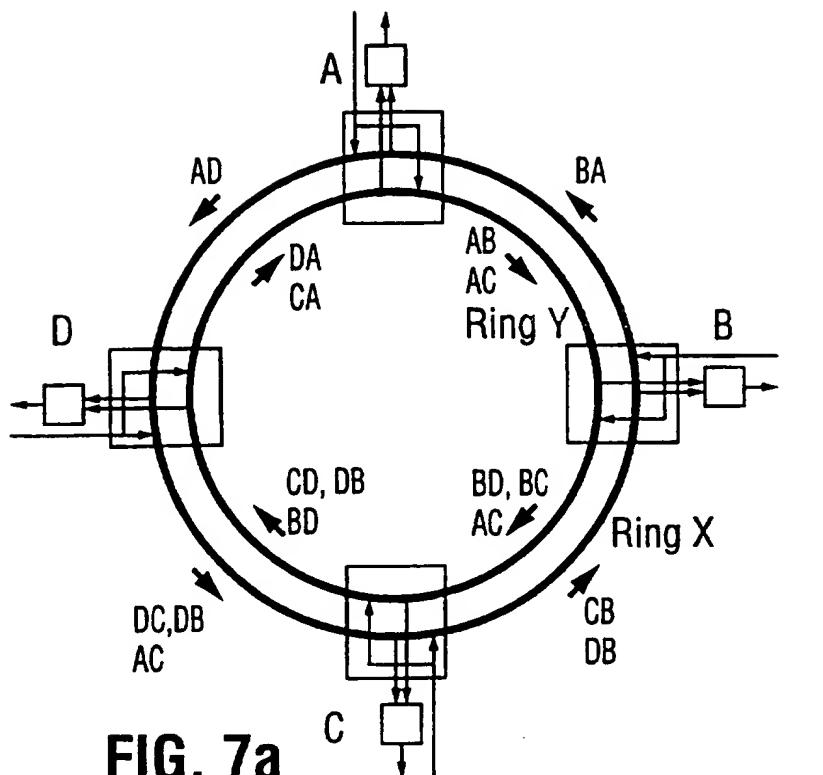


FIG. 7a

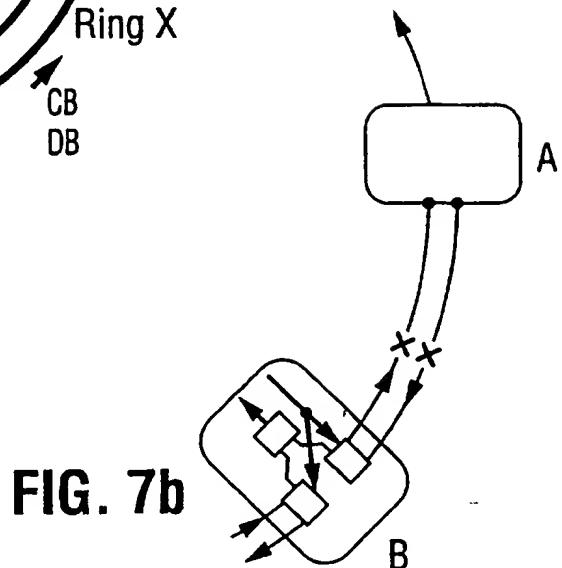


FIG. 7b

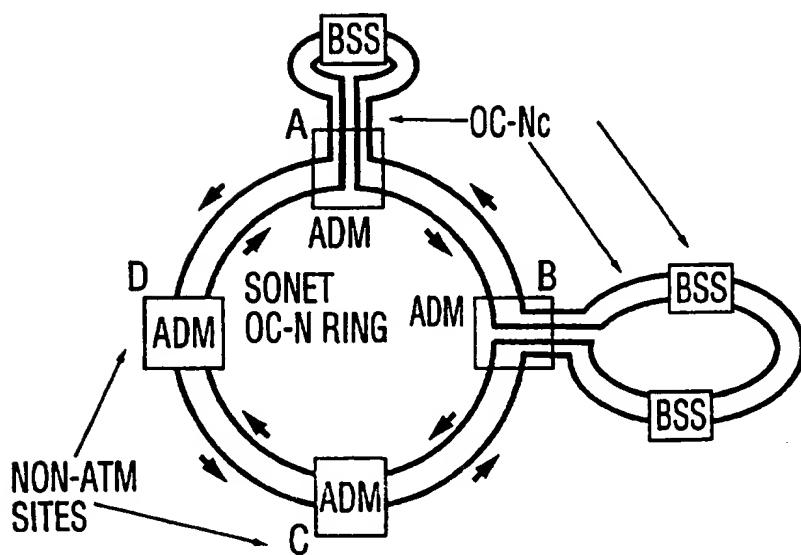


FIG. 8

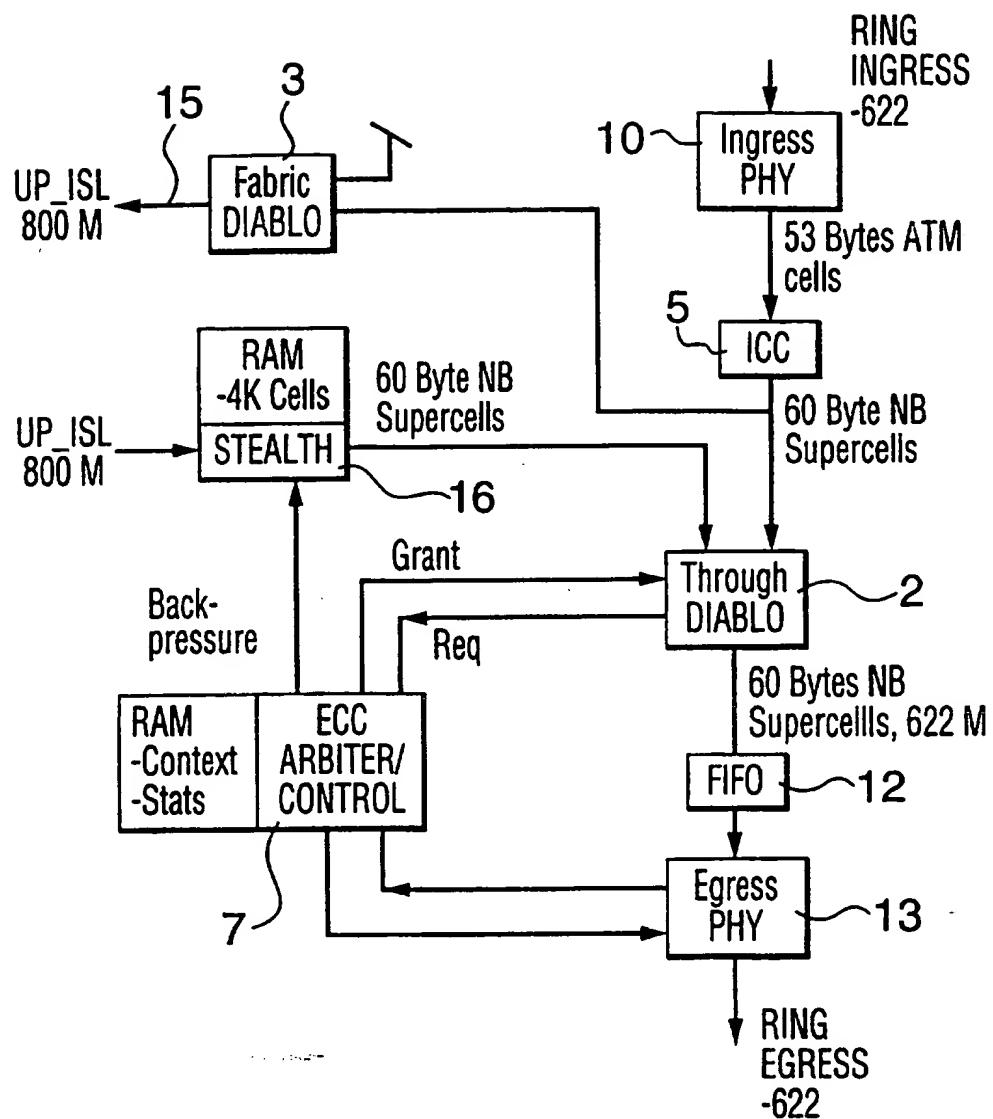


FIG. 9

INTERNATIONAL SEARCH REPORT

International Application No

PCT/CA 97/00052

A. CLASSIFICATION OF SUBJECT MATTER

IPC 6 H04Q11/04 H04L12/43

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 6 H04Q H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	<p>COMMUNICATION FOR GLOBAL USERS, INCLUDING A COMMUNICATIONS THEORY MINI CONFERENCE ORLANDO, DEC. 6 - 9, 1992, vol. 1 - 2 - 03, 6 December 1992, INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS, pages 834-840, XP000357681 TSONG-HO WU ET AL: "A BROADBAND VIRTUAL PATH SONET/ATM SELF-HEALING RING ARCHITECTURE AND ITS ECONOMIC FEASIBILITY STUDY" *section 2* see page 837, column 1, line 7 - line 45 * section 3.1 *</p> <p>---</p> <p>-/-</p>	1-3,5-9, 12,14,15

 Further documents are listed in the continuation of box C. Patent family members are listed in annex.

* Special categories of cited documents :

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *E* earlier document but published on or after the international filing date
- *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- *O* document referring to an oral disclosure, use, exhibition or other means
- *P* document published prior to the international filing date but later than the priority date claimed

T later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

X document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

Y document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

& document member of the same patent family

3

Date of the actual completion of the international search

21 May 1997

Date of mailing of the international search report

09.06.97

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+ 31-70) 340-2040, Tx. 31 651 epo nl.
Fax: (+ 31-70) 340-3016

Authorized officer

Gregori, S

INTERNATIONAL SEARCH REPORT

Inte ~~rnal~~ Application No
PCT/CA 97/00052

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	<p>COMMUNICATION FOR GLOBAL USERS, INCLUDING A COMMUNICATIONS THEORY MINI CONFERENCE ORLANDO, DEC. 6 - 9, 1992, vol. 1 - 2 - 03, 6 December 1992, INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS, pages 639-643, XP000357855 YOSHIO KAJIYAMA ET AL: "ATM SELF-HEALING RING" see page 639, column 1, line 19 - line 25 see page 639, column 2, line 7 - page 640, column 2, line 3 see page 640, column 2, line 29 - page 642, column 1, line 15 ---</p>	1,3-8, 14-18
X	<p>IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS, vol. 12, no. 1, 1 January 1994, pages 171-177, XP000493898 YOSHIO KAJIYAMA ET AL: "AN ATM VP-BASED SELF-HEALING RING" see page 171, column 1, line 18 - column 2, line 16 see page 172, column 2, line 25 - line 32 see page 174, column 2, line 15 - page 175, column 2, line 14 -----</p>	1,8-11, 13